④

AD-A221 601

Stephen Casner
Karen Seo
Winston Edmond
Claudio Topolcic

University
of Southern
California

N–Way Conferencing with Packet Video

DTIC
ELECTE
MAY 16 1990
S E D

90 05 15 056

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION	1b. RESTRICTIVE MARKINGS
Unclassified	

2a. SECURITY CLASSIFICATION AUTHORITY	3. DISTRIBUTION / AVAILABILITY OF REPORT
2b. DECLASSIFICATION / DOWNGRADING SCHEDULE	This document is approved for public release; distribution is unlimited.

4. PERFORMING ORGANIZATION REPORT NUMBER(S)	5. MONITORING ORGANIZATION REPORT NUMBER(S)
ISI/RS-90-252	————

6a. NAME OF PERFORMING ORGANIZATION	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION
USC/Information Sciences Institute		Defense Supply Service Washington

6c. ADDRESS (City, State, and ZIP Code)	7b. ADDRESS (City, State, and ZIP Code)
4676 Admiralty Way Marina del Rey, CA 90292-6695	Room 1D-245 The Pentagon Washington, DC 20310

8a. NAME OF FUNDING / SPONSORING ORGANIZATION	8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER
DARPA		MDA903-87-C-0719 MDA903-89-C-0020

8c. ADDRESS (City, State, and ZIP Code)	10. SOURCE OF FUNDING NUMBERS			
Defense Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, VA 22209	PROGRAM ELEMENT NO. ———	PROJECT NO. ———	TASK NO. ———	WORK UNIT ACCESSION NO. ———

11. TITLE (Include Security Classification)

N-Way Conferencing with Packet Video (Unclassified)

12. PERSONAL AUTHOR(S) Casner, S.; Seo, K.; Edmond, W.; Topolcic, C.

13a. TYPE OF REPORT	13b. TIME COVERED	14. DATE OF REPORT (Year, Month, Day)	15. PAGE COUNT
Research Report	FROM _____ TO _____	1990, April	10

16. SUPPLEMENTARY NOTATION

Reprinted from the proceedings of *The Third International Workshop on Packet Video*, held 22-23 March 1990 in Morristown, New Jersey.

17.	COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	bandwidth reservation, Distributed Queue Dual Bus (DQDB) protocol, multicasting, multimedia conferencing, packet video
09	02		

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

Packet switching technology promises to allow improvement of video quality by efficiently supporting variable-rate video coding. Its inherent multiplexing of multiple streams also allows more efficient multi-destination delivery for N-way conferencing. However, most commercial video codecs are designed to work with circuits, not packets, in part because these benefits are accompanied by some problems. This paper describes a packet video system implementation in which commercial codecs were adapted to exploit the benefits of packet switching while addressing the problems as follows: 1) clock synchronization was obviated by asynchronous operation; 2) delay was reduced by bandwidth reservation and fast packet forwarding; and 3) packet loss was reduced by bandwidth reservation and forward error correction. An overview of the system is followed by sections addressing each of the problems and benefits, plus future directions for expansion of the system.

20. DISTRIBUTION / AVAILABILITY OF ABSTRACT	21. ABSTRACT SECURITY CLASSIFICATION
☒ UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT. ☐ DTIC USERS	Unclassified

22a. NAME OF RESPONSIBLE INDIVIDUAL	22b. TELEPHONE (Include Area Code)	22c. OFFICE SYMBOL
Victor Brown Sheila Coyazo	213/822-1511	

DD FORM 1473, 84 MAR 83 APR edition may be used until exhausted.   SECURITY CLASSIFICATION OF THIS PAGE
All other editions are obsolete.

Stephen Casner
Karen Seo
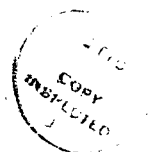Winston Edmond
Claudio Topolcic

University
of Southern
California

N-Way Conferencing with Packet Video

Reprinted from the proceedings of
The Third International Workshop on Packet Video,
held 22-23 March 1990 in Morristown, New Jersey.

*INFORMATION
SCIENCES
INSTITUTE*

N-Way Conferencing with Packet Video

S. Casner, K. Seo*, W. Edmond*, C. Topolcic*

USC/Information Sciences Institute
4676 Admiralty Way, Marina del Rey, CA 90292

*BBN Systems and Technologies Corporation
10 Moulton Street, Cambridge, MA 02238

ABSTRACT

Packet switching technology promises to allow improvement of video quality by efficiently supporting variable-rate video coding. Its inherent multiplexing of multiple streams also allows more efficient multi-destination delivery for N-way conferencing. However, most commercial video codecs are designed to work with circuits, not packets, in part because these benefits are accompanied by some problems. This paper describes a packet video system implementation in which commercial codecs were adapted to exploit the benefits of packet switching while addressing the problems as follows:

1. Clock synchronization was obviated by asynchronous operation;

2. Delay was reduced by bandwidth reservation and fast packet forwarding;

3. Packet loss was reduced by bandwidth reservation and forward error correction.

An overview of the system is followed by sections addressing each of the problems and benefits, plus future directions for expansion of the system.

System Description

The Multimedia Conferencing project, a collaborative effort between ISI and BBN STC under DARPA sponsorship, has developed an experimental system for real-time, multisite conferences [6]. All media (voice, video and shared workspace) are communicated via packet protocols.

The purpose of this system and the underlying network is to provide a platform for the research on high-speed networking protocols and applications: connection-oriented as well as connectionless service, broadcast and multicast service, and real-time conferencing. However, the system is also used regularly for teleconference meetings by sponsors and researchers on various projects.

Conference rooms are installed near San Francisco, Los Angeles, Washington and Boston. A conference may include all four sites, with motion video images from each site displayed simultaneously in quadrants of the video screen. Audio from all sites is mixed for playback so all may talk at once if they wish. The BBN MMCONF system [2] provides the shared workspace, a set of windows that appear identically on a workstation screen at each site. These windows may display text, graphics, bitmap images and other media for presentations and collaborative editing.

The network services used by this conferencing system are supplied by a combination of a backbone network and gateways (routers). The backbone, the Terrestrial Wideband Network, is a wide area network that is part of the initial phase of DARPA's Defense Research Internet (DRI). It provides a linear, trans-continental backbone built on T1 trunks. The Wideband packet switching nodes (WPS) and gateways are based on BBN

Butterfly Multiprocessor hardware. At user sites, local area networks and conferencing equipment are connected to the backbone packet switches via gateways and T1 tail circuits. Figure 1 shows the locations of network nodes and conference sites. Not shown are ten other user sites.
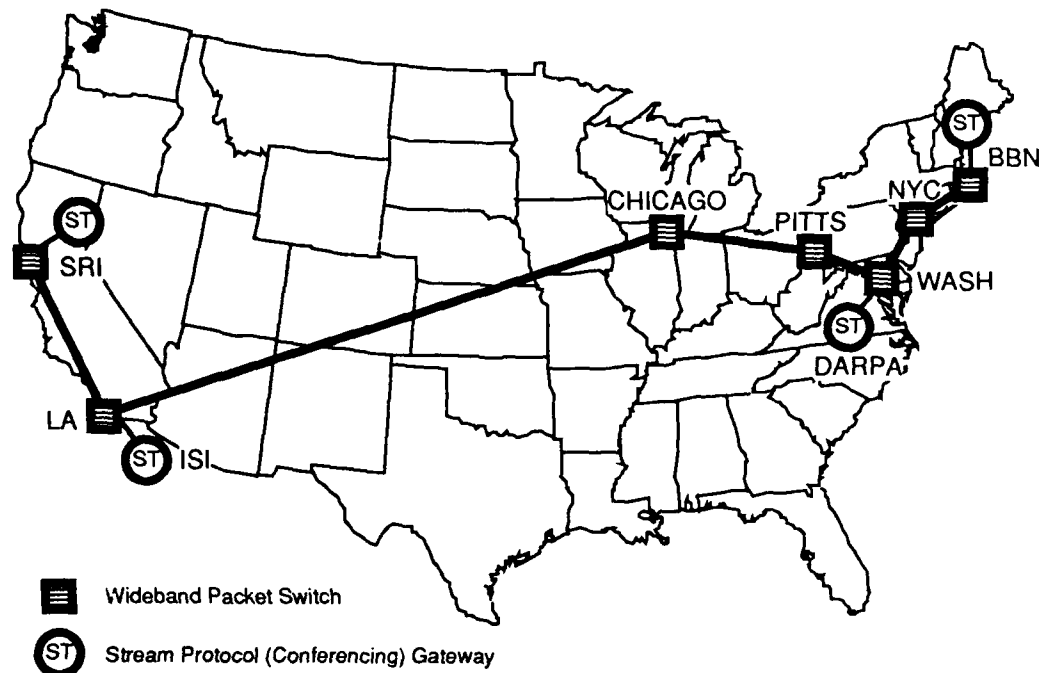

Figure 1: Terrestrial Wideband Network nodes and conference sites

This network manages bandwidth using the BBN Dual Bus Protocol (DBP), a link-level protocol that is a type of Distributed Queue Dual Bus (DQDB) protocol. It is similar to the IEEE 802.6 Metropolitan Area Network (MAN) protocol [4], but with features that support wide area networking and applications such as multimedia conferencing and distributed simulation.

The DBP uses a reservation mechanism that provides access fairness for each WPS. There are two equivalent buses, one in each direction. These buses are slotted, and each bus is used to request slots on the opposite bus. A distributed queuing system is created by using counters in each WPS to match free slots with slot requests. This queuing provides FIFO access to the network. Unlike 802.6, this protocol can also support re-use of packet slots. When the data in a slot reaches its last destination node and does not need to be forwarded, the slot is marked free and can be re-used for sending other traffic further down the line. Additional features are bandwidth preallocation and multicast delivery, described in later sections.

Video and voice traffic is supported by the Stream Protocol (ST) [1,7]. This protocol is at the same level as the DoD Internet Protocol (IP) [3,5] for datagrams, but ST involves an explicit setup phase prior to the data transfer phase. During the setup phase, the application specifies its communication requirements to the ST gateways, which in turn select a route and reserve the necessary bandwidth and other resources in the gateways and the network.
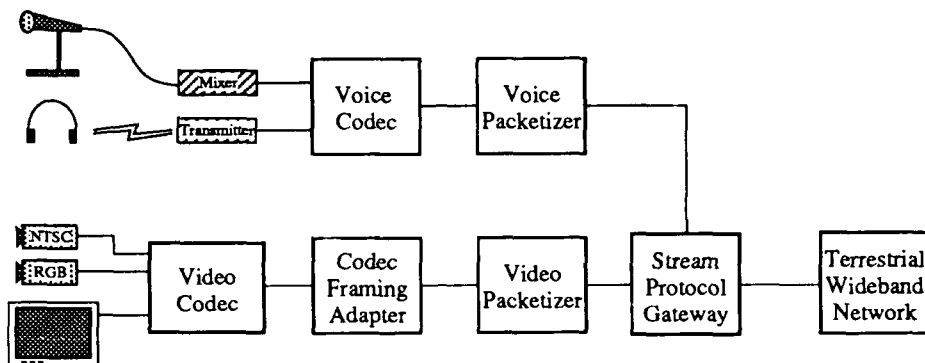
Figure 2: Packet video system components

Problem One: Clock Synchronization

The primary problem posed in packet switching of video is the lack of a clock signal synchronized end-to-end, since circuit-oriented codecs usually expect to receive a continuous stream of bits. A sophisticated solution is to reconstruct a clock at the receiver using phase-locked loop (PLL) techniques based on the depth of the incoming packet queue, but this clock may not be stable enough due to the variance in packet arrival times. A simpler technique can be used if the codec will allow two things:

- a receive clock rate higher than the transmit clock rate, and

- a way to indicate to the receiver when no data is available.

Data bits are transmitted continuously by the codec according to the locally-generated transmit clock, but data bits are delivered to the receiving codec in bursts as packets are received. The transmitter can't overrun the receiver because the receive clock rate is higher, and underrun is not a problem if "no data" can be indicated. No buffer is required to accomodate packet arrival rate jitter, so no delay is added. The image update interval at the receiver may vary slightly, but with low-rate codecs this is generally not noticeable.

This technique has been implemented at ISI for three different commercial circuit-oriented codecs operating in the 56-384 Kb/s range. Using a PC coprocessor card with multiple high-speed serial ports, the proprietary serial communications protocols of the three codecs are followed to extract the native data blocks from each codec. The data blocks are then encapsulated for further processing by the video packetizing software. In the reverse direction, the encapsulation is removed and data blocks are delivered to the codec with the appropriate idle bit pattern repeated between blocks.

Problem Two: Delay

It is important to minimize end-to-end delay because it impedes responsive interaction among participants in a teleconference. Conventional packet store-and-forwarding involves per-packet processing and routing, and buffering of the entire packet at every intermediate node. The DBP minimizes forwarding delay by simplifying forwarding decisions at intermediate nodes in a linear chain. Those packet header fields that change as a packet is forwarded occur first in the header and require only simple processing. The first field is an eight-bit "terminus" field set by the source node. A node can determine whether or not to forward the packet in the time it takes to receive the terminus bits from the trunk and to compare them against the site's

identifier. If there is no match, meaning that the packet doesn't terminate at the site, then the packet can be forwarded along the trunk with only a few byte-times delay. In this way most processing occurs at the entry point and minimal processing and buffering are needed at subsequent nodes along the trunk until the exit point.

In moderately to heavily loaded datagram networks, congestion and queuing can cause high and variable delivery delays, or even packet losses. To avoid these problems, the DBP allows applications to pre-reserve bandwidth by setting up "streams". The application's packets are carried in slots that have been pre-allocated for its stream so that the packets do not have to go through the normal datagram reservation mechanism. This is accomplished by having the network set aside K slots out of each frame of N slots, where K varies according to the bandwidth pre-reserved by the application, and N is a network-wide parameter. Setting aside bandwidth for an application both minimizes delay and reduces its variance because slots are assigned at regular intervals.

Slots assigned to a stream are allocated only between the source and destination nodes. The same slots may be used to carry other stream or datagram traffic between nodes before the source, and between nodes after the destination.

Problem Three: Packet Loss

Packet loss is caused by two factors: queue overflow and transmission errors. Since the ST protocol reserves the necessary network bandwidth and resources to support the video and voice coding rate, the possibility of queue overflow is virtually eliminated.

Since this network is implemented on cross country fiber optic trunks, the transmission error rate is very low. Forward error correction (FEC) could be implemented for the network links to reduce the transmission error rate even further. Errors in the data portion of a packet do not cause a packet loss, so long as the header information is intact, because the network and the ST protocol allow such packets to be delivered to the receiving application. Then FEC implemented in the video codec can correct the data errors.

Disruption of the image due to packet loss is minimized in codecs designed to process incoming frames independently so a complete image refresh is not required. Only one of the three codecs we have adapted has this feature. However, for the other codecs, the low packet loss rate observed in practice still means refreshes are relatively infrequent and therefore not disruptive.

Benefit: Variable Bandwidth

Having addressed these problems of packet switching, we can now exploit its advantages. The first is increased bandwidth efficiency for those video codecs, including one of the three we adapted, that implement constant-quality, variable-data-rate coding.

It may appear that reserving a stream of K slots out of N in a frame is identical to the establishment of a circuit using traditional time-division multiplexing, but there is an important difference. The DBP allows each of the K frames to be used for datagram traffic in the absence of waiting stream traffic. This works because the packet headers allow the two kinds of traffic to be distinguished at the destination. Thus, for variable-rate coding, the peak rate is reserved but capacity between the average and peak rates is not wasted, whereas a circuit would dedicate the peak bandwidth continuously.

Benefit: Multicasting for N-Way Conferencing

A second advantage is improved bandwidth efficiency for N-way conferencing. The DBP allows hosts to set up dynamic multicast groups. Hosts transmit only one copy of each packet to the network, addressed to the group, then the network performs the packet replication and multicast delivery. By letting the network replicate the packet as required, a packet-switched system can take advantage of any overlap in communication paths between group members to reduce the total traffic as compared to point-to-point circuits. Multicast delivery allows linear growth with the number of sites while point-to-point delivery results in quadratic growth, as shown in Figure 3.
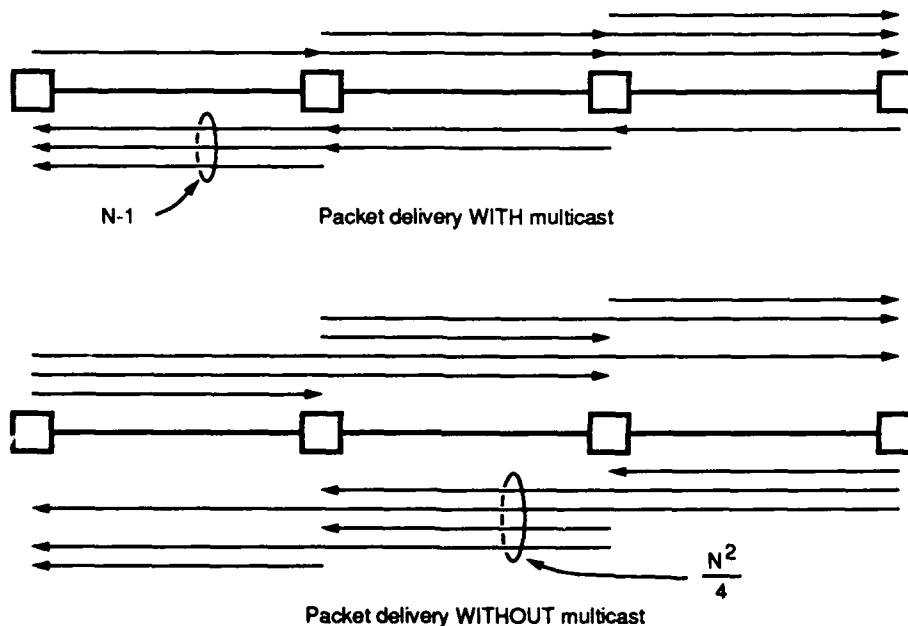

Figure 3: Peak link loading with and without multicast for N=4 sites

Simultaneous Display of Multiple Sites

In an N-way conference, each site receives packets from all the other sites. The packet streams are multiplexed together over a single network connection to the video system, but each packet can be identified by its source address. One of the codecs we use has been modified at ISI to process the multiplexed packet streams from up to four sites and display the images in separate quadrants of the screen for simultaneous viewing. Packets from all the sites can be fed to the codec over a single connection, and need not arrive in round-robin order, because the labels on the packets allow firmware to identify each source. To do the same task with circuits would require multiple connections to the codec or byte-level demultiplexing hardware in the codec.

For codecs that can display only one image at a time, it is still useful for each site to receive packets from all the others. Each site can make an independent selection of the site to be viewed, and deliver only packets from that site to the codec.

Future Directions

We hope to expand this marriage of commercial video codecs and packet networks to bring teleconferencing to more places. T1-based IP datagram networks now link many educational and research institutions. Even though most of these networks do not support resource reservation, those that are lightly loaded and have a small number of nodes may still provide adequate performance for packet video. And, as other types of networks with resource reservation services are developed and added to the internetwork, the ST gateway could easily be modified to access those services.

Another area of this work that we would like to expand is the DBP-based network technology itself. This technology could be scaled to allow handling of more than one trunk per link and to support higher speed trunks, e.g., T3. Networks with more complex topologies could be built using store and forward nodes to interconnect multiple independent DBP chains at the points where they intersect. In this scenario, a packet would gain the fast forwarding advantages of the DBP when traveling on the DBP chains, and would only incur significant processing and buffering delay at the store and forward nodes.

As video coding standards mature, video coding functions may be implemented in VLSI and incorporated directly into workstations along with the audio coding functions that are already appearing. Workstations also incorporate high-speed local-area net connections that could be used for packet video. These features, in combination with gateway and wide-area network technology similar to what we have developed, should make workstation-to-workstation conferencing feasible on a large scale.

Acknowledgements

References

[1] Forgie, J.W., "IEN 119: ST -- A Proposed Internet Stream Protocol" (unpublished memorandum), MIT Lincoln Laboratory (Sept 1979).

[2] Forsdick, H.C., "Explorations into Real-time Multimedia Conferencing", *Proceedings 2nd International Symposium on Computer Message Systems*, pp.299-315 (Sept 1985).

[3] Leiner, B., Cole, R., Postel, J., Mills, D., "The DARPA Internet Protocol Suite", *Proceedings INFOCOM '85, Computers and Communications Integration: The Confluence at Mid-decade* (Mar 1985).

[4] Newman, R.M., Budrikis, Z.L., Hullett, J.L., "The QPSX MAN", *IEEE Communications*, pp.20-28, (Apr 1988).

[5] Postel, J.B., "Internet Protocol", RFC 791, USC/Information Sciences Institute, Marina del Rey, CA (Sept 1981).

[6] Schooler, E.M., Casner, S.L. "A Packet-switched Multimedia Conferencing System", *ACM SIGOIS Bulletin*, pp.12-22 (Jan 1989).

[7] Topolcic, C., Park, P., Draft of "Proposed Changes to the Experimental Internet Stream Protocol (ST)", BBN Laboratories, Cambridge, MA (Apr 1987).